arXiv:1406.1423v1 [cs.DB] 5 Jun 2014

A ToolBox for Conservative XML Schema Evolution and Document Adaptation

Joshua Amavi, Jacques Chabin, Mirian Halfeld Ferrari, Pierre Réty

Univ. Orléans, INSA Centre Val de Loire, LIFO EA 4022, FR-45067 Orléans, France
{joshua.amavi, jacques.chabin, mirian, pierre.rety}@univ-orleans.fr

Abstract. This paper proposes a set of tools to help dealing with XML database evolution. It aims at establishing a multi-system environment where a global integrated system works in harmony with some local original ones, allowing data translation in both directions and, thus, activities on both levels. To deal with schemas, we propose an algorithm that computes a mapping capable of obtaining a global schema which is a conservative extension of original local schemas. The role of the obtained mapping is then twofold: it ensures schema evolution, via composition and inversion, and it guides the construction of a document translator, allowing automatic data adaptation *w.r.t.* type evolution. This paper applies, extends and put together some of our previous contributions.

1 Introduction

The construction of new applications aiming at integrating data from different sources while still allowing the use of original local systems is not an easy task. The idea here is to establish a multi-system environment composed by a global central system which is a *conservative* evolution of local ones, capable of processing changes that can then be transmitted to local systems. The communication should be possible in both directions: local-to-global and global-to-local. The goal is to allow independent local services to continue working on their own data, with their own tools while permitting diagnosis and changes based on a general and complete view of all services. This scenario requires tools for dealing with type evolution and document adaptation. It can be useful as a temporary configuration, deferring complete integration until local systems are ready, or as a flexible architecture adopted by the enterprise.

In this context, we suppose that $S_1, \ldots, S_n$ are local systems which deal with sets of XML documents $X_1, \ldots, X_n$, respectively, and that inter-operate with a global, integrated system S. Each set X_i conforms to schema or type constraints $\mathcal{D}_i$, while $\mathcal{D}$ is an extended type (of S) that accepts any local document from $\mathcal{D}_i$. We assume that the global system S may evolve to S', accepting more documents or rejecting some original ones. Our goal is to propose tools allowing automatic type transformation accompanied by automatic document translation.

The contribution of this paper can be summarized as follows:

- We apply, extend and put together some of our previous work dealing with XML constraints in the context of database integration or evolution.

- In [4] we proposed *ExtSchemaGenerator*, an algorithm for generating a new type which was the closest conservative evolution of some given types. Here, we extend that work by generating *mappings* that indicate how to transform original schemas into the extended one (and vice-versa), via edit operations. These mappings are then used to ensure different forms of type evolution.
- In [2] we developed *XMLCorrector* to correct XML documents *w.r.t.* types. We use it here inside a document translator which is guided by a schema mapping.

The following motivating example illustrates the advantages of our proposal. We then offer an overview of our realisations and goals.

Motivating Example. We consider the hospital data maintained by three services: a service which has information about patients and their treatments, another service that is responsible for bills and one service that keeps contact to insurance companies and tells whether a treatment is covered by the insurance of a patient. Figure 1 shows a summarized version of the DTD of each service. Notice that we omit the definition of elements whose type is PCDATA.

```
Patient and Treatment Service
<!ELEMENT hospital  (info*)>                    <!ELEMENT visitInfo (trId,date)>
<!ELEMENT info      (patient|treatment)>        <!ELEMENT treatment (trId,tname,procedure)>
<!ELEMENT patient   (SSN,pname,visitInfo*)>     <!ELEMENT procedure (treatment*)>

Insurance coverage Service                      Bill Service
<!ELEMENT hospital (info*)>                     <!ELEMENT hospital (info*)>
<!ELEMENT info     (cover|policy)>              <!ELEMENT info     (bill)>
<!ELEMENT cover    (SSN,plname)>                <!ELEMENT bill     (SSN,item*,date)>
<!ELEMENT policy   (plname,trId*)>              <!ELEMENT item     (trId,price)>
```

Fig. 1. DTD of the services of a hospital

Without interfering with these local services, possibly demanding to keep their own local systems, the hospital direction may want to have a global view of all services to process reports and statistics as, for instance, the percentage of insurance companies covering radiotherapy or myopia surgery; the number of patient paying treatments by their own, etc. The global system may receive information directly: a doctor having access to the global schema may introduce information about a new treatment (the price he fixed for this treatment) together with the first patients he is going to treat. Moreover, by analysing all global data, the direction may decide to change its politics. For example, it should decide to introduce a discount for patients not been covered at all, provoking a schema modification (or evolution).

This flexibility can be reached by permitting some basic actions. Firstly, the construction of a global conservative schema capable of accepting any local document together with new documents submitted directed to the global schema. Notice that the local data does not need to be translated, since it is valid *w.r.t.* the global schema. Secondly, the translation of documents from the global to a local system, allowing updates made on the global level to be passed to the local level. Finally, the evolution of the global schema keeping available a mapping to

translate local source schema to the new global one. Similar actions are allowed to deal with local schema evolution.

Figure 2 shows the global DTD resulting from the method we have proposed in [4]. In this paper, we introduce an algorithm that builds mappings from the original DTD to the global one (built according to the ideas in [4]). This mapping is expressed in terms of a sequence of edit operations. Knowing the mapping that transforms a local schema to the global one, it is straightforward to obtain its inverse. The inverse mapping will be the basis for translating documents from a global schema to a local one.

```
<!ELEMENT hospital  (info*)>
<!ELEMENT info      ((patient|treatment)|(cover|policy)|bill)>
<!ELEMENT patient   (SSN,pname,visitInfo*)>      <!ELEMENT cover    (SSN,plname)>
<!ELEMENT visitInfo (trId,date)>                 <!ELEMENT policy   (plname,trId*)>
<!ELEMENT treatment (trId,tname,procedure)>      <!ELEMENT bill     (SSN,item*,date)>
<!ELEMENT procedure (treatment*)>                <!ELEMENT item     (trId,price)>
```

Fig. 2. Global DTD for the Hospital

In Figure 3(a) we find a document valid *w.r.t.* the billing local schema. Notice that it is also valid *w.r.t.* the global schema of Figure 2. Figure 3(b) shows an XML document concerning patients and bills. This document is valid *w.r.t.* the global schema but not valid *w.r.t.* to any local schemas. Translating the document of Figure 3(b) into a document respecting the patient schema we obtain the document of Figure 3(c). The given translation is guided by the schema mapping from the global schema of Figure 2 to the local patient schema of Figure 1.

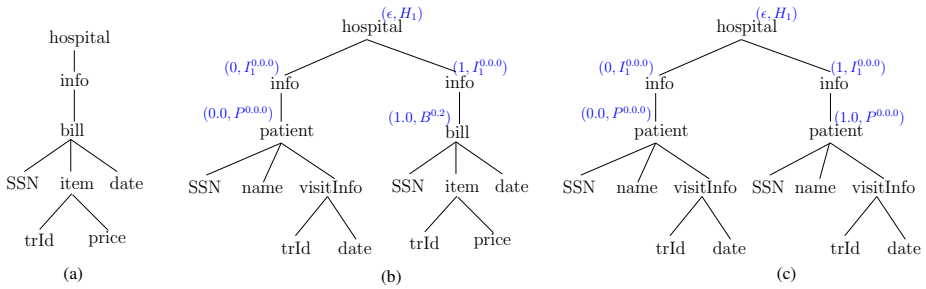

Fig. 3. (a) An XML tree valid *w.r.t.* the billing local schema. (b) An XML tree valid *w.r.t.* the global schema of Figure 2. (c) Tree resulting from the translation of (b) into the patient local schema of Figure 1. Trees (b) and (c) are annotated (c.f. Section 4.2)

Tools for Supporting Schema Evolution. We propose a set of tools to help dealing with XML database evolution. Our goal is to implement a platform where all our proposed tools will be available. Below we describe some important modules of our *ToolBox*, distinguishing those that have been proposed and implemented previously and those that we introduce in the current paper.

Tools from previous work:[1]

- *ExtSchemaGenerator*([4]) extends a given schema G, seen as a regular tree grammar, into a new grammar G' respecting the following property: the language generated by G' is the smallest set of unranked trees that contains the language generated by G and the grammar G' is a Local Tree Grammar (LTG) or a Single-Type Tree Grammar (STTG).
- *XMLCorrector*([2]) corrects an XML document *w.r.t.* schema constraints expressed as a DTD (or an LTG). The corrector reads the entire XML document (or tree) t in order to propose solutions. *XMLCorrector* finds all solutions within a given threshold th.

The above tools are the fundamental bricks for the new ones proposed in this paper. The ideas introduced by *ExtSchemaGenerator* are followed in order to build our mapping generator while *XMLCorrector* is called by our translation module to correct parts of an XML document. The new proposed methods are essential tools to allow schema evolution and compatible document translation.

Tools introduced in the current work:

- *MappingGen*: We propose an algorithm that applies the ideas of [4] to generate a mapping from one schema G, seen as a regular tree grammar, to an extended schema G' which will be an LTG. The resulting schema mapping m is a sequence of operations on grammar rules that indicates, step by step, how to transform G into G' following the approach in [4]. Given a mapping m we can easily compute its inverse m^{-1} or compose it to other mappings; allowing schemas to evolve.
- *XTraM*: Based on a given mapping m (from schema S to T), we propose a method to translate an XML document (or tree) t, valid *w.r.t.* S into a document t' valid *w.r.t.* T. The edit distance between t and t' is no higher than a given positive threshold th. Moreover, t' is the closest tree to t, obtained by changing t according to the schema modifications imposed by m. For each edit operation on S, to obtain T, we analyse what should be the corresponding update on document t. When this update violates validity, we use *XMLCorrector* to propose corrections to the subtree involved in the update.

2 Background

An XML document is an *unranked tree*, defined in the usual way as a mapping t from a set of positions $Pos(t)$ to an alphabet Σ. The set of the trees over Σ is denoted by T_Σ. For $v \in Pos(t)$, $t(v)$ is the label of t at the position v. *Positions* are sequences of integers in $\mathbb{N}^*$ and $Pos(t)$ satisfies: $\forall u, i, j\,(j \geq 0, u.j \in Pos(t), 0 \leq i \leq j) \Rightarrow u.i \in Pos(t)$ (char "." denotes the concatenation). The *size* of t (denoted $|t|$) is the cardinal of $Pos(t)$. As usual, ϵ denotes the empty sequence of integers, i.e. the root position and t, t' will denote trees.

[1] *ExtSchemaGenerator* is available on `http://www.univ-orleans.fr/lifo/Members/rety/logiciels/RTGalgorithms.html`. *XMLCorrector* is available on `http://www.info.univ-tours.fr/~savary/English/xmlcorrector.html`.

Given a tree t, we denote by $t|_p$ the subtree whose root is at position $p \in Pos(t)$, *i.e.* $Pos(t|_p) = \{s \mid p.s \in Pos(t)\}$ and for each $s \in Pos(t|_p)$ we have $t|_p(s) = t(p.s)$. Now, let $p \in Pos(t)$ and t' be a tree, we note $t[p \leftarrow t']$ as the tree that results of substituting the subtree of t at position p by t'.

Definition 1 (Regular Tree Grammar, derivation). A *regular tree grammar* (RTG) is a 4-tuple $G = (N, \Sigma, S, P)$, where: N is a finite set of *non-terminal symbols*; Σ is a finite set of *terminal symbols*; S is a set of *start symbols*, where $S \subseteq N$ and P is a finite set of *production rules* of the form $X \to a\,[R]$, where $X \in N$, $a \in \Sigma$, and R is a regular expression over N. We say that, for a production rule, X is the left-hand side, $a\,[R]$ is the right-hand side, and R is the content model.
For an RTG $G = (N, \Sigma, S, P)$, we say that a tree t built on $N \cup \Sigma$ derives (in one step) into t' iff (i) there exists a position p of t such that $t|_p = A \in N$ and a production rule $A \to a\,[R]$ in P, and (ii) $t' = t[p \leftarrow a(w)]$ where $w \in L(R)$ ($L(R)$ is the set of words of non-terminals generated by R). We write $t \to_{[p, A \to a\,[R]]} t'$. More generally, a derivation (in several steps) is a (possibly empty) sequence of one-step derivations. We write $t \to^*_G t'$.
The *language* $L(G)$ generated by G is the set of trees containing only terminal symbols, defined by: $L(G) = \{t \mid \exists A \in S,\ A \to^*_G t\}$. □

Remark: As usual, in this paper, our algorithms start from grammars in reduced form and (as in [9]) in normal form. A *regular tree grammar* (RTG) is said to be in **reduced form** if (i) every non-terminal is reachable from a start symbol, and (ii) every non-terminal generates at least one tree containing only terminal symbols. A regular tree grammar (RTG) is said to be in **normal form** if distinct production rules have distinct left-hand-sides. □

Among RTG we are particularly interested in local tree grammars which have the same expressive power as DTD[2]. We recall the definition from [10]:

Definition 2 (Local Tree Grammar). Two non-terminals A and B (of the same grammar G) are said to be *competing with each other* if $A \neq B$ and G contains production rules of the form $A \to a[R]$ and $B \to a[R']$ (i.e. A and B generate the same terminal symbol). A *local tree grammar* (LTG) is a regular tree grammar that does not have competing non-terminals[3]. A *local tree language* (LTL) is a language that can be generated by at least one LTG. □

3 Schema Evolution

3.1 Conservative XML Type Extension (*ExtSchemaGenerator*)

In [4] we find conservative evolution algorithms that compute a local or single-type grammar which extends minimally a given original regular grammar. That

[2] Note that converting an LTG into normal form produces an LTG as well.
[3] In contrast, a *single-type tree grammar* (STTG) is an RTG in normal form, where (i) for each production rule, non terminals in its regular expression do not compete with each other, and (ii) start symbols do not compete with each other.

paper proves the correctness and the minimality of the generated grammars. In the current paper we will only deal with the generation of LTG. We follow the idea of *ExtSchemaGenerator* which is very simple when dealing with the generation of an LTG from an RTG: replace each pair of competing non-terminals by a new non-terminal, until there are no more competing non-terminals. The regular expression of a new non-terminal rule is the disjunction of the regular expressions associated to competing non-terminals.

Let us consider the example of Section 1 where we have three hospital services, each one having its own LTG (or DTD) as schema. Figure 4 shows the RTG obtained by the union of the production rules of all these three grammars while Figure 5 shows the resulting LTG. The obtained LTG is an extension of the original RTG since it generates *all* trees generated by the original RTG and possibly others as well (refer to example of Figure 3). Clearly, the obtained grammar is also an extension of each hospital service grammar.

$H_1 \rightarrow hospital[I_1^*]$
$I_1 \rightarrow info[P \mid T]$
$P \rightarrow patient[S \cdot N \cdot V^*]$
$V \rightarrow visitInfo[Id \cdot D]$
$T \rightarrow treatment[Id \cdot TN \cdot PR]$
$PR \rightarrow procedure[T^*]$

$H_2 \rightarrow hospital[I_2^*]$
$I_2 \rightarrow info[C \mid Pol]$
$C \rightarrow cover[S \cdot PN]$
$Pol \rightarrow policy[PN \cdot Id^*]$

$H_3 \rightarrow hospital[I_3^*]$
$I_3 \rightarrow info[B]$
$B \rightarrow bill[S \cdot It^* \cdot D]$
$It \rightarrow item[Id \cdot PZ]$

Fig. 4. RTG obtained from the union of production rules of grammars.

$H_1 \rightarrow hospital[I_1^* \mid I_1^* \mid I_1^*]$
$I_1 \rightarrow info[(P \mid T) \mid (C \mid Pol) \mid B]$
$P \rightarrow patient[S \cdot N \cdot V^*]$
$V \rightarrow visitInfo[Id \cdot D]$
$T \rightarrow treatment[Id \ \cdot \ TN \ \cdot \ PR]$
$PR \rightarrow procedure[T^*]$
$I_1 \rightarrow info[(P \mid T) \mid (C \mid Pol) \mid B]$
$Pol \rightarrow policy[PN \cdot Id^*]$
$B \rightarrow bill[S \cdot It^* \cdot D]$
$It \rightarrow item[Id \cdot PZ]$

Fig. 5. LTG obtained by algorithm in [4] from the RTG of Figure 4.

3.2 Schema Mappings

In the context of schema evolution, we say that a *source* schema (or grammar) evolves to a *target* schema. A *schema mapping* is specified by an operation list, denoted as an *edit script*, that should be performed on source schema in order to obtain the target schema. In this paper, we propose an algorithm that generates a mapping to translate an RTG G into an LTG G', following the lines of [4] . Our mapping is composed by a sequence of *edit operations* that should be applied on the rules of grammar G in order to obtain G'. Before defining all our edit operations we formally introduce the notions of edit script and schema mapping.

In the following definition, let *ed* be an *edit operation* defined on RTG G. We denote by $ed(G)$ the RTG obtained by applying *ed* on G. Each edit operation is associated with a cost that can be fixed according to the user's priority. Thus, the cost of an edit script is the sum of the costs of the edit operations composing it.

Definition 3 (Edit Script and Edit Script Cost). An *edit script* $m = \langle ed_1, ed_2, \ldots ed_n \rangle$ is a sequence of *edit operations* ed_k where $1 \leq k \leq n$. Let G be an RTG, an *edit script* $m = \langle ed_1, ed_2, \ldots ed_n \rangle$ is defined on G if and only if there exists a sequence of RTG $G_0, G_1, \ldots, G_n$ such that: (i) $G_0 = G$ and (ii) $\forall\, 1 \leq k \leq n$, ed_k is defined on G_{k-1} and $ed_k(G_{k-1}) = G_k$. Hence, we have $m(G) = G_n$. The empty *edit script* is denoted $\langle\rangle$. The cost of an edit script m is defined as $cost(m) = \Sigma_{i=1}^{n}(cost(ed_i))$. □

Definition 4 (Schema Mapping). A schema mapping is a triple $\mathcal{M} = (S, T, m)$, where S is the source schema, T is the target schema, and m is an *edit script* that transforms S into T (*i.e.*, $m(S) = T$). We say that $\mathcal{M}$ is syntactically *specified by*, or, *expressed* by m. □

3.3 Edit operations

In this section, we define edit operations on an RTG $G = (N, \Sigma, S, P)$. The idea is, firstly, to represent production rules as trees. Then, the problem of changing one RTG into another is treated as a tree editing problem.

Tree Representation for Production Rules. Let $X \rightarrow a\,[R]$ be a production rule. We denote by $reg(X)$ the regular expression R associated with the non-terminal X, and by $term(X)$ the terminal symbol a for the non-terminal X. Note that $reg(X)$ and $term(X)$ are defined in a deterministic way since we always suppose that grammars are in normal form, therefore X occurs only once as the left-hand-side of a production rule. We treat the regular expression R as an unranked tree denoted t_R. The set of non-terminal symbols occurring in R is denoted by $nt(R)$. Formally, t_R is recursively defined as follows:

- if $R = \epsilon$ then t_R is a single node labeled by ϵ.
- if $R = A$ where $A \in N$ then t_R is a single node labeled by A.
- if $R = R_1.\cdots.R_n$ then $t_R = .(t_{R_1}, \cdots, t_{R_n})$ *i.e.* t_R is tree such that the root is '.' with the subtrees $t_{R_1}, \cdots, t_{R_n}$.
- if $R = R_1|\cdots|R_n$ then $t_R = |(t_{R_1}, \cdots, t_{R_n})$ *i.e.* t_R is tree such that the root is '|' with the subtrees $t_{R_1}, \cdots, t_{R_n}$.
- if $R = R_1^*$ then $t_R = *(t_{R_1})$.
- if $R = (R_1)$ then $t_R = t_{R_1}$.

We represent the right-hand side of a production rule $X \rightarrow a\,[R]$ as a tree denoted t_X^r such that $t_X^r = a(t_R)$. The root of t_X^r is the terminal a which has only one subtree t_R. We have that $t_X^r|_0 = t_R$. For example, in Figure 6, the tree on the top left corner is t_I^r with $I \rightarrow info[T.(Y \mid Co)]$.

Definition 5 (Well Formed Tree). A tree t representing the right-hand side of a production rule is well formed iff the following conditions are verified:
(i) the root is a terminal symbol, *i.e.* $t(\epsilon) \in \Sigma$, and has exactly one child;
(ii) the leaves nodes are in $N \cup \{\epsilon\}$ and
(iii) the internal nodes are in the set $\{|, ., *\}$ such that: if an internal node is in $\{*\}$ then the internal node has exactly one child; otherwise if an internal node is in $\{|, .\}$ then the internal node has at least one child. □

Elementary Edit Operations. We define elementary edit operations by using rewriting. Given a set of variables X, a rewrite rule (written $l \to r$) is a pair of terms over $\{|, ., *\} \cup N \cup X$, assuming that variables have no children. A hedge is a (possibly empty) sequence of trees, like $[t_0, \ldots, t_n]$. Let h be a hedge, $|h|$ denotes the number of trees in h. For example, if $h = [t_0, \ldots, t_n]$, then $|h| = n + 1$. A substitution σ is a mapping of finite domain from X into the set of hedges, whose application is extended homomorphically to trees. Let t, t' be trees, t rewrites into t' at position u, by the rewrite rule $l \to r$, and with the substitution σ (written $t \to_{[u, l \to r, \sigma]} t'$) if $t|_u = \sigma(l)$ and $t' = t[u \leftarrow \sigma(r)]$. For example, given the rule $f(x, y) \to g(y, x)$ (x and y are variables), then $f(a, b, c)$ rewrites into $g(c, a, b)$ (with substitution $x/[a, b]$, $y/[c]$), and also into $g(b, c, a)$ (with $x/[a]$, $y/[b, c]$).

In the following definition, terms are always rewritten at position u, with substitution σ, provided the condition (if any) is satisfied. We only mention the rewrite rule, which is not always the same.

Definition 6 (Elementary Edit Operations). Given an RTG $G = (N, \Sigma, S, P)$ in normal form, an *elementary edit operation ed* is a partial function that transforms G into a new RTG G'. The *elementary edit operation ed* can be applied on G only if *ed* is defined on G. We distinguish four types of *elementary edit operations* on RTG:

1. Edit operations to modify the set of start symbols S
   - `set_startelm`(A): adds the non-terminal A to S where $A \in N$.
   - `unset_startelm`(A): deletes the non-terminal A from S where $A \in N$.
2. Edit operations to modify non-terminal or terminal symbols in a content model
   - `ins_elm`$(X, A, u.i)$: (cf. Figure 6(ed_1)) applies the rewrite rule
     $op(x, y) \longrightarrow op(x, A, y)$
     on t^r_X at position u where $X \in N$, $A \in N \cup \{\epsilon\}$, $|\sigma(x)| = i$ and $op \in \{|, .\}$.
   - `del_elm`$(X, A, u.i)$: (cf. Figure 6(ed_2)) applies the rewrite rule
     $op(x, A, z) \longrightarrow op(x, z)$
     on t^r_X at position u where $X \in N$, $|\sigma(x)| = i$, $A \in N \cup \{\epsilon\}$, $|\sigma(x)| + |\sigma(z)| \geq 1$ and $op \in \{|, .\}$.
   - `rel_root`(X, a, b): (cf. Figure 6(ed_3)) applies the rewrite rule
     $a(x) \longrightarrow b(x)$
     on t^r_X at position ϵ where $X \in N$, $a, b \in \Sigma$ and $|\sigma(x)| = 1$.
   - `rel_elm`(X, A, B, u): (cf. Figure 6(ed_4)) applies the rewrite rule $A \longrightarrow B$ on t^r_X at position u where $X \in N$, $A, B \in N \cup \{\epsilon\}$.
3. Edit operations to modify operator symbols in a content model
   - `ins_opr`$(X, opr, u.i, n)$: (cf. Figure 6(ed_5)) applies the rewrite rule

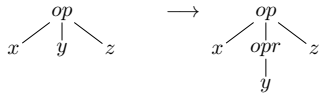

on t^r_X at position u where $X \in N$, $n \geq 1$, $op \in \{|, ., *\} \cup \Sigma$, $|\sigma(x)| = i$, $|\sigma(y)| = n$ and if $n = 1$ then $opr \in \{|, ., *\}$ otherwise $opr \in \{|, .\}$.

- `del_opr`$(X, opr, u.i, n)$: (cf. Figure 6(ed_6)) applies the rewrite rule

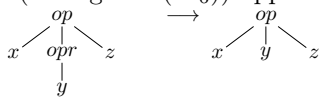

on t^r_X at position u where $X \in N$, $op \in \{|, ., *\} \cup \Sigma$, $opr \in \{|, ., *\}$, $|\sigma(x)| = i$, and $|\sigma(y)| = n$. If $op \in \{*\} \cup \Sigma$ then $|\sigma(y)| = 1$ and $|\sigma(x)| + |\sigma(z)| = 0$.

- `rel_opr`(X, op, opr, u): (cf. Figure 6(ed_7)) applies the rewrite rule

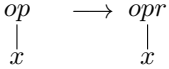

on t^r_X at position u where $X \in N$, $op, opr \in \{|, ., *\}$, $|\sigma(x)| \neq 0$ and if $opr = *$ then $|\sigma(x)| = 1$.

4. Edit operations to modify the set of production rules P
   - `ins_rule`(A, a): adds the new production rule $A \to a\,[\epsilon]$ to P and the non-terminal A to S, where $A \notin N$.
   - `del_rule`(A, a): deletes the production rule associated with A from P, where $A \in N$ and $reg(A) = \epsilon$. If $A \in S$ then A is also deleted from S.

After each edit operation, the sets Σ and N are automatically updated to contain all and only the terminal (resp. non-terminal) symbols appearing in P. □

Proposition 1. *An edit operation applied on an RTG G results in an RTG G' that is also in normal and reduced form.* □

Proposition 2. *Let G and G' be two RTG. There exist an edit script, composed only by operations of Definition 6, that transforms G into G'.* □

Non-Elementary Edit Operations. For readability and cost estimation, we define short-cut operations, *i.e.*, operations seen as a one-block operation but equivalent to a sequence of *elementary edit operations*. In this paper, we introduce just those that are used in our algorithm:

- `ins_tree`$(X, R, u.i)$ (and, respectively, `del_tree`$(X, R, u.i)$): consists of inserting (respect. deleting) a subtree at a given position in the right-hand side of a rule. This operation is similar to `ins_elm`(respect. `del_elm`) but instead of adding (respect. deleting) a node in t^r_X, it adds (respect. deletes) the subtree t_R.
- `ins_treerule`(A, a, R) (and, respectively, `del_treerule`(A, a, R)): adds (respect. deletes) the new (respect. existing) productive rule $A \to a\,[R]$ to P and the non-terminal A to S, where $A \notin N$ (respect. where $A \in N$) and $nt(R) \subseteq N$. We can transform this operation in a sequence that adds $A \to a\,[\epsilon]$ to P and then changes the regular expression ϵ into R (respect., that changes t_R into ϵ and then deletes the rule $A \to a\,[\epsilon]$ from P).

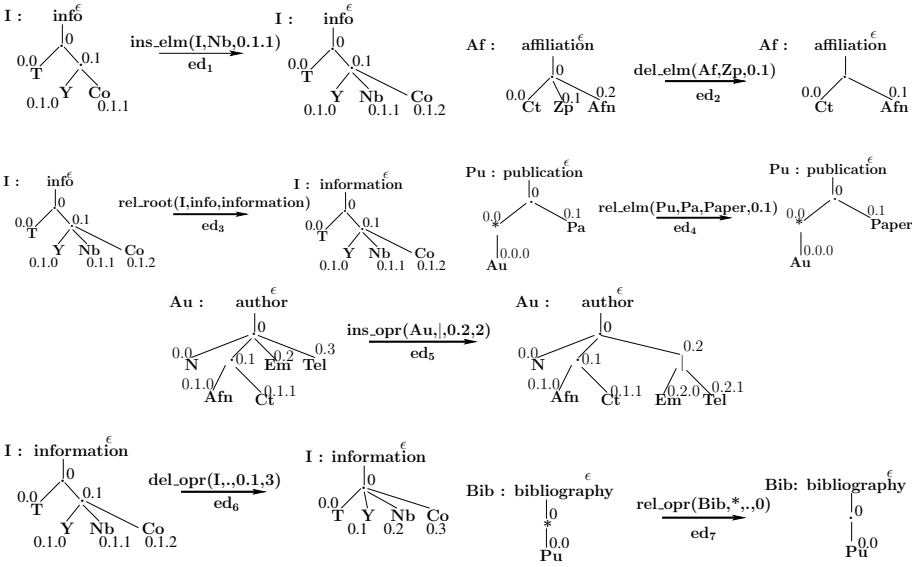

Fig. 6. Example of *elementary edit operations*.

Now, for each *edit operation ed*, we define a non-negative and application-dependent cost. On the one hand, we assume that operations that do not change the language generated by the RTG G on which they were applied, are 0-cost. Their goal is just to simplify a given regular expression. For instance, $\texttt{del_opr}(X, opr, u.i)$ where $t^r_X(u) = t^r_X(u.i) = opr$ and $\texttt{del_opr}(X, opr, u.i)$ where $t^r_X(u.i) \in \{|, .\}$ and $t^r_X(u.i)$ has exactly one child, are 0-cost operations. On the other hand, we suppose that an *elementary edit operation* (Definition 6) costs 1, while a *non-elementary edit operation* costs 5.

3.4 Generating a Schema Mapping (*MappingGen*)

Algorithm 1 generates a mapping that converts an RTG in an LTG by following the ideas in [4], explained in Section 3.1. This algorithm starts by determining a set of competing non-terminals EC_a (lines 2-3). Then we can take arbitrarily in EC_a, one of these non-terminals (say X_0) to represent all others, *i.e.*, when merging rules of competing terminals, one non-terminal name is chosen to represent the result of the merge (line 4). Recall that edit operations always deal with a production rule in its tree-like format. The new production rule of X_0 is built in two steps. We add an OR operation as the parent of its original regular expression $reg(X_0)$ (line 5) and then we insert all regular expressions associated with its competing non-terminals as siblings of $reg(X_0)$ (line 7). In line 8 we just replace, in all production rules, non-terminals in EC_a by X_0. Original rules of non-terminals in EC_a are deleted (line 10) after, possibly, adjusting start symbols (line 9).

Algorithm 1 A mapping for transforming an RTG into an LTG

Input: A Regular Tree Grammar $G = (NT, \Sigma, S, P)$
Output: An *edit script* m between G and the LTG G' such that $L(G) \subseteq L(G')$

$m := \langle\rangle$
for each terminal symbol $a \in \Sigma$ **do**
  $EC_a = \{X_0, \ldots, X_k\}$ is a set of competing non-terminals where $term(X_i) = a$
  Non-terminal X_0 is choosed to represent $X_0, \ldots, X_k$
  Add $\texttt{ins_opr}(X_0, |, 0, 1)$ to m
  for each non-terminal $X_i \in \{X_1, \ldots, X_k\}$ **do**
    Add $\texttt{ins_tree}(X_0, reg(X_i), 0.i)$ to m
    Add $\texttt{rel_elm}(Y, X_i, X_0, u)$ to m, for all u where u is the position of X_i in the rule $Y \rightarrow b\,[R] \in P$
    Add $\texttt{set_startelm}(X_0)$ to m where $X_0 \notin S$ and $X_i \in S$
    Add $\texttt{del_treerule}(X_i, a, reg(X_i))$ to m
  end for
end for
return m

Consider the RTG of Figure 4. Algorithm 1 returns the following mapping m: $\langle \texttt{ins_opr}(H_1, |, 0, 1)$, $\texttt{ins_tree}(H_1, reg(H_2), 0.1)$, $\texttt{del_treerule}(H_2, hospital, reg(H_2))$, $\texttt{ins_tree}(H_1, reg(H_3), 0.2)$, $\texttt{del_treerule}(H_3, hospital, reg(H_3))$, $\texttt{ins_opr}(I_1, |, 0, 1)$, $\texttt{ins_tree}(I_1, reg(I_2), 0.1)$, $\texttt{rel_elm}(H_1, I_2, I_1, 0.1.0)$, $\texttt{del_treerule}(I_2, info, reg(I_2))$, $\texttt{ins_tree}(I_1, reg(I_3), 0.2)$, $\texttt{rel_elm}(H_1, I_3, I_1, 0.2.0)$, $\texttt{del_treerule}(I_3, info, reg(I_3))\rangle$. When m is applied on the RTG of Figure 4, the LTG of Figure 5 is obtained.

Proposition 3. *Let m be the mapping obtained by Algorithm 1 from an RTG G. The language $L(m(G))$ is the least LTL that contains $L(G)$. Moreover, the grammar $m(G)$ equals the one obtained by ExtSchemaGenerator.* □

3.5 Going Further with Mappings to support Schema Evolution

In [5], it was shown how two fundamental operators on schema mappings, namely composition and inversion, can be used to address the mapping adaptation problem in the context of schema evolution. Given $\mathcal{M}_1 = (S, T, m_1)$, a mapping between XML schemas S and T, when S or T evolve, $\mathcal{M}_1$ shall be adapted. By using composition and inversion operators, one can avoid mapping re-computation. The idea is illustrated in Figure 7. Firstly, suppose the target schema evolves to T', and that this evolution is modeled by mapping $\mathcal{M}_2$. Composing $\mathcal{M}_1$ and $\mathcal{M}_2$, denoted by $\mathcal{M}_1 \circ \mathcal{M}_2$, is an operation that has the same effect as applying first $\mathcal{M}_1$ and then $\mathcal{M}_2$. For example, given an RTG G, Algorithm 1 obtains $\mathcal{M}_1$ and we can find an LTG G_1. If G_1 evolves into G_2 using $\mathcal{M}_2$, the translation of the original G into G_2 is obtained just by computing $\mathcal{M}_1 \circ \mathcal{M}_2$. Now, suppose that the source schema evolves to a new source schema S', modeled by mapping $\mathcal{M}_3$. To obtain schema T' from S', the composition of $\mathcal{M}_3$ with $\mathcal{M}_1 \circ \mathcal{M}_2$ is not possible, since $\mathcal{M}_3$ and $\mathcal{M}_1 \circ \mathcal{M}_2$ are not consecutive. To apply composition we need, first, to compute the inversion of $\mathcal{M}_3$, denoted $\mathcal{M}_3^{-1}$, which "undoes" the

effect of $\mathcal{M}_3$. Once we obtain a suitable $\mathcal{M}_3^{-1}$, we can then apply the composition operator to produce $\mathcal{M}_3^{-1} \circ \mathcal{M}_1 \circ \mathcal{M}_2$. The resulting schema mapping is now from S' to T'. We now precise the notions of composition and inversion in our context.

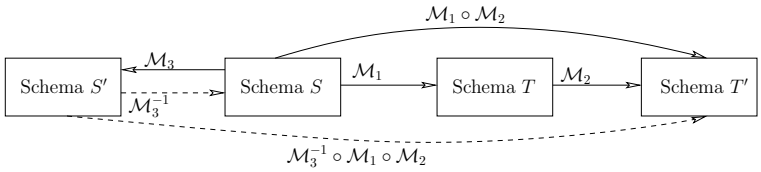

Fig. 7. Application of composition and inversion in schema evolution.

Definition 7 (Mapping composition and inversion). Given two mappings $\mathcal{M}_1 = (S, T, m_1)$ and $\mathcal{M}_2 = (T, V, m_2)$, the composition of $\mathcal{M}_1$ and $\mathcal{M}_2$ is the mapping $\mathcal{M}_1 \circ \mathcal{M}_2 = (S, V, m_1 \,.\, m_2)$. If $m_1 = \langle ed_1, \cdots, ed_n \rangle$, then the inverse of mapping $\mathcal{M}_1$ is the mapping $\mathcal{M}_1^{-1} = (T, S, m_1^{-1})$ where $m_1^{-1} = \langle ed_n^{-1}, \cdots, ed_1^{-1} \rangle$ and $ed_k^{-1} (1 \leq k \leq n)$ is defined in Table 1. □

$\texttt{set_startelm}(A) \rightleftharpoons \texttt{unset_startelm}(A)$	$\texttt{ins_opr}(X, p, u, n) \rightleftharpoons \texttt{del_opr}(X, p, u, n)$
$\texttt{ins_elm}(X, A, u) \rightleftharpoons \texttt{del_elm}(X, A, u)$	$\texttt{rel_opr}(X, p, q, u) \rightleftharpoons \texttt{rel_opr}(X, q, p, u)$
$\texttt{rel_root}(X, a, b) \rightleftharpoons \texttt{rel_root}(X, b, a)$	$\texttt{ins_rule}(A, a) \rightleftharpoons \texttt{del_rule}(A, a)$
$\texttt{rel_elm}(X, A, B, u) \rightleftharpoons \texttt{rel_elm}(X, B, A, u)$	$\texttt{ins_tree}(X, R, u) \rightleftharpoons \texttt{del_tree}(X, R, u)$
$\texttt{ins_treerule}(A, a, R) \rightleftharpoons \texttt{del_treerule}(A, a, R)$	

Table 1. Inverse relationship for edit operations.

To illustrate the inverse operation, consider the mapping generated by Algorithm 1 for the RTG of Figure 4 (Section 3.4). The inverse of this mapping is:$\langle \texttt{ins_treerule}(I_3, info, reg(I_3))$, $\texttt{rel_elm}(H_1, I_1, I_3, 0.2.0)$, $\texttt{del_tree}(I_1, reg(I_3), 0.2)$, $\texttt{ins_treerule}(I_2, info, reg(I_2))$, $\texttt{rel_elm}(H_1, I_1, I_2, 0.1.0)$, $\texttt{del_tree}(I_1, reg(I_2), 0.1)$, $\texttt{del_opr}(I_1, |, 0, 1)$, $\texttt{ins_treerule}(H_3, hospital, reg(H_3))$, $\texttt{del_tree}(H_1, reg(H_3), 0.2)$, $\texttt{ins_treerule}(H_2, hospital, reg(H_2))$, $\texttt{del_tree}(H_1, reg(H_2), 0.1)$, $\texttt{del_opr}(H_1, |, 0, 1)\rangle$. This inverse mapping, applied on the LTG of Figure 5, gives the RTG of Figure 4.

4 Adapting XML Documents to a New Type

4.1 Correcting XML Documents (*XMLCorrector*)

In [2], given a well-formed XML tree t, a schema G and a non negative threshold th, *XMLCorrector* finds every tree t' valid *w.r.t.* G such that the edit distance between t and t' is no higher than th. Contrary to most other approaches, [2] considers the correction as an enumeration problem rather than a decision problem and computes all the possible corrections on t. The algorithm, proved to be correct and complete in [2], consists in fulfilling an edit distance matrix which stores the relevant edit operation sequences allowing to obtain the corrected trees. The theoretical exponential complexity of *XMLCorrector* is related to the

fact that edit sequences and the corresponding corrections are generated and that the correction set is complete.

In this paper, contrary to [2], we do not consider all the possible corrections on t. The correction of XML documents is guided by a given mapping. For each edit operation on S, to obtain T, we analyse what should be the corresponding update on document t. When this update violates validity, we use *XMLCorrector* to propose corrections to the subtree involved in the update.

4.2 Document Translation guided by Mapping (*XTraM*)

This section outlines our data translation method which is guided by a schema mapping. Our method consists in performing a list of changes on XML documents, in accordance with the edit operations found in the mapping. For example, adding or deleting a regular expression in a rule under the operator '.' is a mapping operation that provokes, respectively, the insertion or the deletion of a subtree in an originally valid XML tree (to maintain its validity). Similarly, renaming a non-terminal A by B, provokes the substitution of the subtree generated by A into the subtree generated by B. When local correction on XML subtrees are needed, *XMLCorrector* is used to ensure document validity.

Consider an XML tree t valid *w.r.t.* schema S and a mapping m from S to T. Our method can be summarized in two steps:

1. Since t belongs to the language $L(S)$, it is possible to associate a non-terminal A with each tree node position p generated by this non-terminal. We analyse t, detect each non-terminal and annotate it with its corresponding position u in the used production rule. This annotation respects the format (p, A^u). For example, in Figure 3(b), we notice that the tree node *bill* is generated by the non-terminal B whose position in $t^r_{I_1}$ is 0.2, noted as $(1.0, B^{0.2})$.
2. Each edit operation ed in m activates a set of modifications on t. When ed transforms a grammar into a new grammar containing the previous one, the set of modifications is empty. Otherwise, our method consists in traversing t (marked as in step 1) in order to find the tree positions which may be affected due to ed. Modifications on t are defined according to each edit operation and are not detailed here due to the lack of space. Obviously, if no position is affected, t does not change.

The inverse mapping (Section 3.5) guides changes on tree t_1 (Figure 3(b)).

1. Nodes in t_1 are annotated in blue, Figure 3(b).
2. `ins_treerule`$(I_3, info, B)$ implies zero change on t_1. By inserting a new production rule in our grammar in Figure 5, we obtain a new grammar which contains the previous one. The new grammar generates also t_1.
3. `rel_elm`$(H_1, I_1, I_3, 0.2.0)$ requires each I_1 s.t. $t^r_{H_1}(0.2.0) = I_1$ to be renamed I_3. As in t_1 (Figure 3(b)), there is no annotation where $I_1^{0.2.0}$ is a child of annotation (ϵ, H_1), no changes are performed on t_1.
4. `del_tree`$(I_1, B, 0.2)$ requires to delete B s.t. $t^r_{I_1}(0.2) = B$. As in Figure 3(b) annotation $(1.0, B^{0.2})$ (node labeled *bill*) is a child of annotation $(1, I_1^{0.0.0})$ changes should be performed on t_1. Since we delete B from the expression

$(P \mid T) \mid (C \mid Pol) \mid B$ we must replace the subtree generated by B by a subtree generated by $(P \mid T) \mid (C \mid Pol)$ for preserving the validity of t_1 *w.r.t.* the new grammar. For doing that, we launch *XMLCorrector* on subtree at position 1.0 in t_1 and the expression $(P \mid T) \mid (C \mid Pol)$ that, in this case, computes only the subtree *patient(SSN,name,visitInfo(trId,date))* with a (minimal) cost of 5. Let t_2 be the tree with the new subtree (Figure 3(c)).

5. All other operations in the mapping imply zero change on t_2 due to the same reasons as stated in (2) or (3). As expected, the result is in Figure 3(c).

5 Related Work and Concluding Remarks

Much other work deals with schema evolution. In [5] authors show how inversion and composition models schema evolution, by expressing schema constraints in a logical formalism. Second order logic is needed to express some mapping compositions. This approach is the basis for proposals in [1, 7, 13] dealing with XML schema evolution. We believe that the use of edit operations makes our approach simpler than theirs and gets on well with our previous work concerning XML document correction. Other proposals, such as those in [6, 8, 11, 3, 12], use edit operations. ELaX (Evolution Language for XML-Schema) in [11] and Exup [3] are a domain-specific language that proposes to handle modifications on XSD and to express such modifications formally. Contrary to us, approaches in [8, 12] only consider LTG evolution. In [6] we find a proposal that is closer to ours, dealing with RTG evolution. An important originality of our approach is the *automatic* generation of a *conservative extension* of an RTG into an LTG, following the lines of [4]. The use of a schema mapping to guide document adaptation is also considered in [12, 3]. However in [12], when the original grammar is ambiguous allowing more than one solution, their method fails. Our approach may propose different solutions to be chosen by the user. *XTraM* is guided by a mapping and produces documents with corrections that do not exceed a threshold.

ExtSchemaGenerator([4]) returns a conservative extended grammar. *MappingGen* automatically produces a mapping for this conservative type evolution. Having a mapping allows any evolution (conservative or not) via inversion or composition. *XTraM* uses *XMLCorrector* locally and follows a given mapping to propose XML document adaptations. Thus, all *local* solutions under a given threshold are produced. Our system offers flexibility. In *XTraM* (Section 4.2, step 1), different annotations are possible (indeed, for 1-ambiguous LTG only one annotation is possible, but general RTG allow distinct annotations). Our method can produce all possible document adaptations (*i.e.*, those respecting local thresholds) or let to the user the choice of following just a fixed number of them. The user can also adapt edit operation costs according to his priorities. A prototype, implemented in Java, is been tested. As a first experiment, we have produced an LTG, in $24ms$, by merging the grammars obtained from *dblp* DTD[4] and *HAL XSD*[5]. *MappingGen* returned a 19-operation mapping. Then *XTraM*

[4] `http://dblp.uni-trier.de/xml/dblp.dtd`

[5] `http://import.ccsd.cnrs.fr/xsd/generationAuto.php?instance=hal`

was used to adapt a 52-node document valid *w.r.t.* the computed LTG toward the *HAL* grammar, giving, in this case, 36 solutions in 22.6 *s*. As in this test, all possible translations can be considered, but the user may also interfere in an intermediate step, making choices before the end of the complete computation - guiding and, thus, restricting the number of solutions. We are currently working on a friendly interface to facilitate this intermediate interference.

Our *ToolBox* offers schema evolution mechanisms accompanied by an automatic adaptation of XML documents. Its conservative aspect guarantees great flexibility when a global integrated system co-exists with local ones.

References

1. Amano, S., Libkin, L., Murlak, F.: XML schema mappings. In: Proceedings of the 28th ACM Symposium on Principles of Database Systems. pp. 33–42. PODS '09, ACM, New York, NY, USA (2009)
2. Amavi, J., Bouchou, B., Savary, A.: On correcting XML documents with respect to a schema. The Computer Journal 56(4) (2013)
3. Cavalieri, F., Guerrini, G., Mesiti, M.: Updating XML schemas and associated documents through Exup. In: Proceedings of the IEEE 27th International Conference on Data Engineering. pp. 1320–1323. ICDE '11, IEEE Computer Society, Washington, DC, USA (2011)
4. Chabin, J., Halfeld Ferrari Alves, M., Musicante, M.A., Réty, P.: Conservative type extensions for XML data. T. Large-Scale Data- and Knowledge-Centered Systems, TLDKS Journal 9(LNCS 7980), 65–94 (2013)
5. Fagin, R., Kolaitis, P.G., Popa, L., Tan, W.C.: Schema mapping evolution through composition and inversion. In: Bellahsene, Z., Bonifati, A., Rahm, E. (eds.) Schema Matching and Mapping, pp. 191–222. Springer (2011)
6. Horie, K., Suzuki, N.: Extracting differences between regular tree grammars. In: Proceedings of the 28th Annual ACM Symposium on Applied Computing. pp. 859–864. SAC '13, ACM, New York, NY, USA (2013)
7. Jiang, H., Ho, H., Popa, L., Han, W.S.: Mapping-driven XML transformation. In: Proceedings of the 16th International Conference on World Wide Web. pp. 1063–1072. WWW '07, ACM, New York, NY, USA (2007)
8. Leonardi, E., Hoai, T.T., Bhowmick, S.S., Madria, S.K.: Dtd-diff: A change detection algorithm for dtds. Data Knowledge Engineering 61(2), 384–402 (2007)
9. Mani, M., Lee, D.: XML to Relational Conversion using Theory of Regular Tree Grammars. In: VLDB Workshop on EEXTT. pp. 81–103. Springer (2002)
10. Murata, M., Lee, D., Mani, M., Kawaguchi, K.: Taxonomy of XML schema languages using formal language theory. ACM Trans. Inter. Tech. 5(4), 660–704 (2005)
11. Nösinger, T., Klettke, M., Heuer, A.: XML schema transformations. In: Proceedings of the 24th International Conference on Database and Expert Systems Applications. DEXA '13, vol. 8055, pp. 293–302. Springer Berlin Heidelberg (2013)
12. Suzuki, N., Fukushima, Y.: An XML document transformation algorithm inferred from an edit script between DTDs. In: Proceedings of the 19th Conference on Australasian Database - Volume 75. pp. 175–184. ADC '08, Australian Computer Society, Inc., Darlinghurst, Australia (2007)
13. Yu, C., Popa, L.: Semantic adaptation of schema mappings when schemas evolve. In: Proceedings of the 31st International Conference on Very Large Data Bases. pp. 1006–1017. VLDB '05, VLDB Endowment (2005)